InspireTrace: A Generative AI System for Creative Scaffolding with Source Attribution

Erica Lai[1], Qingyu Huang[2*], Xiaoyue Cheng[2*], Yu Chen[3]

[1]University of California, Irvine
[2]Instilling Goodness Developing Virtue School
[3]San Jose State University

elai14@uci.edu, anna.huang@igdvs.org, hannah.cheng@igdvs.org, yu.chen@sjsu.edu

Abstract

The rapid integration of generative AI into education, industry, and everyday creative practices has introduced significant efficiencies, and yet it has also raised ethical concerns around authenticity, originality, and the erosion of human creativity. Many generative AI systems prioritize speed and content production over user development. Simply delivering synthesized outputs tends to limit opportunities for users to focus on the creative process and creative thinking. Additionally, users might unintentionally plagiarize others' ideas from the synthesized output from AI systems without realizing it. In this paper, we present *InspireTrace*, an early-stage generative AI prototype designed to support creativity through guided inspiration rather than direct solution, intentionally leaving synthesis and ideation to the user. By explicitly attributing source materials, *InspireTrace* aims to reduce plagiarism risks and foster a healthier, more trustworthy creativity ecosystem. We discuss the system's design rationale and its implications from creative practice in human-AI collaboration.

Introduction

The ongoing trend of generative AI integration into education, industry, and everyday lives is transforming the way humans view the world. AI is playing a major role in influencing creativity and innovation. Generative AI has proved to bring myriads of advantages in academia including substantial reduction in researchers' workloads, time-saving mechanisms, the extraction of valuable insights from extensive datasets, and an overall enhancement in the quality of scholarly outputs (Panda & Kaur, 2024). It has also presented new opportunities for organizations to fundamentally rethink their business capabilities and understand new connections between people (e.g., staff, customers, agents of other organizations, and other relevant actors), processes, and technology through AI-enabled automation, engagement, insight, and innovation (Holmström & Carrol 2025). However, this rapid acceleration and integration has also sparked problems regarding authenticity, originality, and creativity.

Despite the significant productivity gains accessed by generative AI, the current industry and the system itself focuses on delivering final products over the developmental process of human creativity. Nonetheless, studies suggest that achieving satisfactory levels of accuracy remains a challenge in many industrial applications with participants expressing concerns about the impact of inaccuracies on trust and user satisfaction (Yu et al. 2025). Hence, shifting the focus of these generative AI systems from efficiency to prioritize user development and assistance may therefore enhance the human creative agency, independent thinking, and creative confidence necessary for long-term innovation (Habib et al. 2024). In addition, research has shown that engaging in creative activities (e.g., composing a song, writing a script) allows individuals to enter flow states which are conducive to one's life satisfaction and psychological well-being (Tan et al. 2021). This enhances emotional stability, reinforces their perception of individual contribution, and minimizes feelings of unease related to potential overreliance on AI.

Generative AI's retrospective nature bound by its training data and pattern recognition causes its lack of traceable co-creativity. Holmström and Carrol (2025) mention through their study that given ChatGPT's reliance on existing data, it makes it more challenging to think outside the box regarding radical innovations (i.e., to think differently, unconventionally, or from a new perspective). The so-called "creative" ideas that generative AI provide, which may look novel to its users, are merely aggregates from existing datasets, and therefore are arguably novel or creative on a higher level.

Additionally, given that nearly all generative AI responses contain paraphrasing and summarizing of what is available on the web, there is a lack of acknowledgement

* These authors contributed equally to this paper

and attribution to the original content creators. Users, therefore, adopt ideas from AI recommendations not necessarily knowing that these responses are fundamentally derived from others' work. This creates a transparency gap that leaves users vulnerable to unintentional plagiarism.

Currently, a widely used application to decrease hallucination in AI is Retrieval-Augmented Generation (RAG), which optimizes the performance of an AI model by connecting it with external knowledge bases, allowing it to present accurate information with source attribution. However, because its primary objective is to deliver concrete and specific answers, it remains insufficient for addressing our broader goal of enhancing human creativity and well-being at an internal, psychological level.

To address the research gap, we designed and developed an early prototype called *InspireTrace* that aims to reimagine a new paradigm of AI for creativity emphasizing on human autonomy, creativity, and intellectual property by keeping principles in mind: 1) provide inspiration to users; 2) leave room for users to come up with their own new ideas; 3) acknowledge the source of origin. Although it applies a retrieval structure similar to RAG, the primary purpose of *InspireTrace* is distinctly different from that of RAG. In contrast, *InspireTrace* is designed to serve as a scaffold for creativity, prioritizing the preservation of the user's originality rather than offering concrete solutions. Another aspect that distinguishes *InspireTrace* from RAG is that *InspireTrace* directly quotes and acknowledges the source of any information it uses. This shift from helping the user find sources to helping the user understand originality could address the ethical concerns mentioned in the abstract. Through the use of restrictive prompts given to the system, it presents ideas and clarifies ownership, thereby enhancing transparency and authenticity. This approach is effective against the risk of plagiarism, fostering a healthier system, and ensuring that intellectual property is fully protected.

Related Work

Current research demonstrates that humans, when assisted by generative AI, can significantly increase productivity in coding, ideation, and written assignments while raising concerns regarding potential disinformation and stagnation of knowledge creation (Zhou & Lee 2024). However, within the interdisciplinary domains of arts and information systems, an essential question remains continuously explored: how creative is generative AI?

While generative AI has demonstrated its potential to boost productivity and positive critique, studies also show that it may be pushing artists toward visual homogeneity and decline in content novelty (Zhou & Lee 2024). This is supported by a study conducted by Stevenson et al. (2022) to measure the creativity of generative AI using the Alternative Use Test (AUT), who found that humans outperform GPT-3 when tasked with creating original creative output. These findings suggest that while generative AI can be an impressive collaborator, it is still lacking the ability to think out-of-the-box which is necessary for solving complex open-ended questions alone known as the trade-off between novelty and coherence [effectiveness] (Cropley 2025).

Generative AI falls under the umbrella of ML and represents an approach where machines can generate new content or data that is similar but not necessarily identical to what they have been trained on (Boussioux et al. 2024). Rather than being able to bring up new ideas, AI is built upon a pattern recognition that is retrospective and ultimately confined by the specific data that it was exposed to and trained on (Boussioux et al. 2024). Hence, LLMs tend to reflect more mainstream ideas due to statistical convergence to the "mean" answer unless directed otherwise (Anderson et al. 2024). This technical feature brings up a significant ethical dilemma, which although may be unintentional, can be overstepping the boundaries of which academia deems original. Because the "new" ideas suggested by generative AI are based on many people's creativity, users often apply these ideas unknowing of its originality. Without mechanisms to trace these ideas back to its origins, the creative process in collaboration with AI is at risk of falling into the traps of plagiarism— 'the use of words and/or ideas from another source, without appropriate attribution' or 'representation of them as one's original work', whether intentionally or unintentionally (Fatemi & Saito 2019).

In a study conducted by Habib et al. (2024) to test the impact of generative AI towards students' creativity, students reflected that they rely on AI as taking "the easy way out" or taking "someone else's work" without credit. Notably, 9 out of 53 students stated that while they thought AI can be helpful in brainstorming, they would prefer not to use AI to assist them. This hesitation from the students arise from the realization that these models build their foundation based on pre-existing human knowledge without the capacity to independently generate novel ideas.

Such concerns reflect broader sentiments within the artistic communities, where AI is perceived as a threat that steals from the works upon which the models are trained and infringes on the copyrights of artists (Zhou & Lee 2024). This produces a transparency gap where the separation between AI-driven prediction and human judgement constantly gets blurrier (Rafner et al. 2023), making it tough for creators to claim the full ownership of their product. Therefore, without a mechanism to properly attribute and synthesize ideas to their origins, creative fields may become saturated with generic content, potentially stifling exploration of new creative frontiers (Zhou & Lee 2024). This ethical dilemma highlights an urgent need for legal scholars to consider the costs and benefits of co-creativity in terms of intellectual property

rights and to develop systems that preserve a human sense of purpose in the age of automation (Rafner et al. 2023).

Ultimately, addressing these ethical and creative concerns requires moving beyond human-AI competition towards human-centered AI and hybrid intelligence (Rafner et al. 2023). This involves adopting "AI-in-the-loop" workflows that keeps humans closely in the loop by utilizing humans' prompt engineering to guide language models in generating creative outputs (Boussioux et al. 2024). In this paradigm, which researchers term "generative synesthesia," the focal skills of creativity shift from mechanical execution to human-led ideation and filtering. To prevent the statistical convergence to the mean answer and the resulting visual homogeneity, researchers advocated a differentiated search approach (Boussioux et al. 2024).

In order to transform the generative process from an opaque summary of the web into a think tank of traceable ideas, we propose a system called *InspireTrace* that supports users to brainstorm with AI building on existing work while explicitly acknowledging original creators. This framework aims to foster mutual learning between humans and machines, ensuring that AI remains a collaborative and ethical creative agent that reduces overreliance on AI and honors original content creators on where the recommendations are based.

InspireTrace is different from traditional search engines in its capability of understanding and analyzing the question and generating responses based on user prompt. Conventional search engines return a large amount of results, requiring users to filter the relevance, availability and reliability. Additionally, the results returned are based on keyword matching, meaning information is displayed as long as key words appear – regardless of whether the content truly answers the user's question. Instead, our system leverages natural language processing to identify underlying user needs and deliver more relevant responses that aim to provide inspiration

Prototype Design and Implementation

Different from prior work, our proposed system is not designed to generate creative content for users directly, but serves the role as an intermediary between users and existing knowledge. We considered the following three design rationales: no generation without evidence, defaulting to citation rather than paraphrasing, and preserving space for user-driven synthesis and decision-making.

In this early version of the prototype, we start by considering a simple scenario: an AI-powered think tank system that provides creative idea recommendations to users based on prior ideas from other creators. We used a database that functions as the sole knowledge source for the system. It consists of a collection of existing ideas, their original text, creator information, and links to primary sources if any. All system generated outputs are strictly constrained to information retrieved from this database, ensuring that responses are grounded in verifiable and traceable references.

We mainly focus on designing prompts used to guide the system output and constrain the AI's capacity for creative reconstruction. Under this constraint, the system output is limited to using only the data provided by the system and is required to respond exclusively with direct quotations from original texts. As a result, the prompt prevents simple summarization, expansion, or hallucinated outputs without any reference or attribution to original content creators.

The written prompt within our code that instructs/programs our generative AI model is structured into three primary sections. First, the system is defined as a research assistant that has no personal opinions and no capabilities beyond matching user requirements with the existing database, extracting quotations, and presenting them to the user. Second, the prompt specifies a fixed response format that the system must follow, which includes: (1) the relevant idea identified in the database, (2) some details of the ideas including referencing from the original text, and (3) the corresponding content creators or source. Third, the prompt requires the system to explain in the output the relevance between the user's input and the retrieved ideas such as variables, units, and nouns found in the knowledge base. In this first prototype, in order to keep the scenario simple, this current version of the system does not allow the output to introduce external content beyond the internal dataset.

In alignment with the system's objectives, users are not provided with definitive or synthesized answers. Instead, the system is designed and defaulted to provide inspiration to users and to support the creative process by supplying relevant source materials for human interpretation, synthesis, and inspiration. After a user submits a question describing their needs, the system retrieves related texts and evaluates whether the database contains sufficient evidence to address the user's request. If adequate references are available, the system selects the most relevant sources by citing original text (if any), original content creator, source links, and a brief contextual explanation on the rationale behind the recommendation. In case of a lack of sufficient references, the system declines to answer and provides a justification for this limitation.

We have developed a web-based application as a prototype that validates the feasibility of the proposed design. The system employs the GPT-4.0 model provided by OpenAI, utilizes an Excel spreadsheet containing a lightweight knowledge base for preliminary testing, and uses Streamlit to build the user interface.

Scenario

To illustrate how this system operates in a real-world setting, we ground the design in an educational use case. We

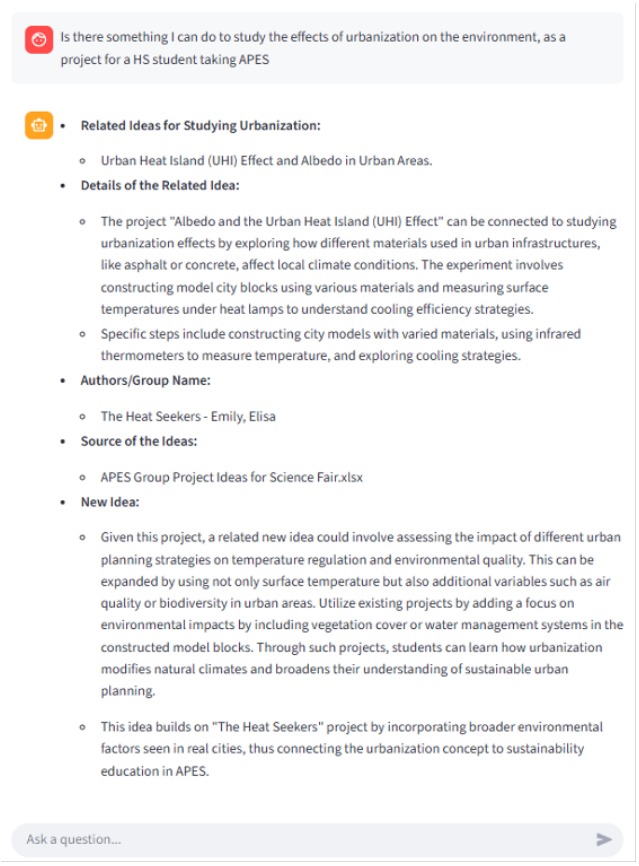

Figure 1: Screenshot of the example scenario using *InspireTrace.*

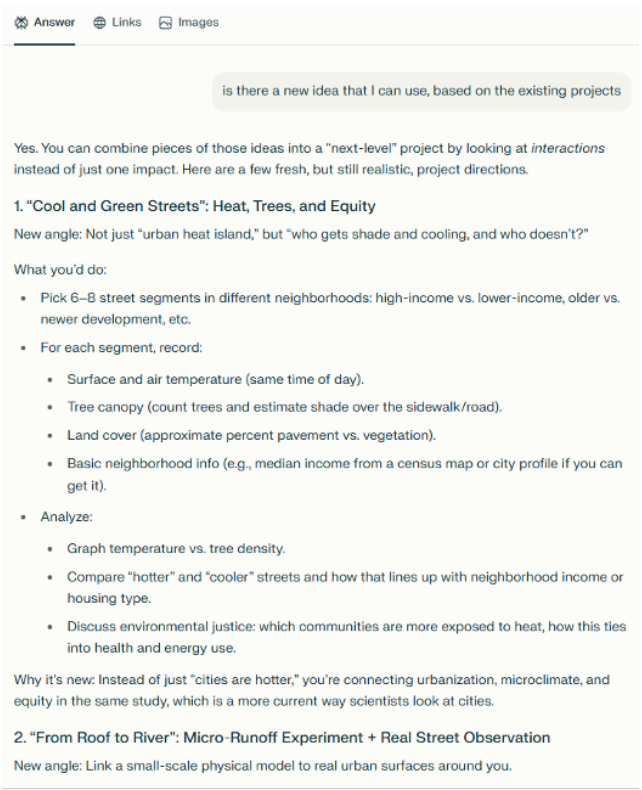

Figure 2: Screenshot of the example scenario using Perplexity AI.

focus on a secondary-school learning context scenario, where students often seek guidance and inspiration from existing work but must also navigate issues of attribution, evidence, and original and independent thinking.

Picture Hana, a 10th grader who's currently taking AP Environmental Science (APES), and she has begun to recognize how class concepts connect to real-world issues such as water and air quality – problems that she observed personally as she grew up in a heavily polluted city. Motivated to explore how urbanization affects the environment, Hana is interested in developing a project idea but is unsure how to begin without simply searching existing work online.

As an initial attempt, Hana turns to the *InspireTrace* system to help, similar to how students often brainstorm with teachers and peers. Rather than directly generating novel project ideas directly, as what commonly used generative AI tools such as ChatGPT could offer, *InspireTrace* is provided with a spreadsheet of prior AP Environmental Science projects, including original project information, contributor information as a curated knowledge base. In response to Hana's question related to ideation, the system identifies relevant existing projects through the GPT-4 model, offers directly cited excerpts, and clearly attributes each idea to its original creators. By presenting prior work as inspiration rather than as synthesized answers and ideas, the system supports Hana in understanding the landscape of existing ideas while preserving space for her own interpretation, ideation and independent project development, while acknowledging the prior work that Hana's work might be inspired by or built upon in the future. See Figure 1 for the output of the above mentioned scenario.

As *InspireTrace* is currently built on top of GPT-4, we cannot guarantee that it is entirely free from hallucinations or inaccuracies in its responses. However, compared to other generative AI systems, *InspireTrace* has a significantly lower likelihood of hallucinating because it is strictly prompted to ground its responses in the provided database and is explicitly permitted to refrain from answering when there is insufficient evidence. Therefore, rather than claiming to eliminate inaccuracies, *InspireTrace* is designed to constrain responses and reduce the frequency of unreliable outputs.

Meanwhile, compared to the response generated by current evidence-based AIs, who tend to analyze and summarize the key points on top of providing the evidence, *InspireTrace* aims to serve as the scaffolding to the new ideas without giving direct solutions/ideas to users. *InspireTrace* allows users to keep up with their critical thinking and analysis skills by leaving the synthesizing part up to the user—

the crucial key to creativity. *InspireTrace* automatically provides ideas for future works that the users may use as reference. By entering a similar prompt to Perplexity AI shown in Figure 2, we asked the system to suggest additional project ideas. The response generated was pedantic, featuring overly precise steps of actions that users can just follow without even thinking or doubting the validity, hence limiting the opportunity for users to think.

Compared to the response generated by many generative AI systems, which typically analyze, summarize key points, and present supporting evidence, *InspireTrace* aims to scaffold for new ideas rather than providing step-by-step instructions. Through intentionally leaving the synthesis and final synthesis process to users, *InspireTrace* encourages the continued development of critical thinking and analytical skills, which are essential for creativity. In addition, *InspireTrace* automatically suggests directions for future work that users may choose to explore.

Future Work

The proposed features and usages of *InspireTrace* acts as the basis to a system that allows for responsible and transparent human-AI interactions that foster a healthy creativity ecosystem. Further developments of the current system may include, but are not limited to, gathering specific data for different contexts, perfecting the AI-user interactions through continuous user output and prompt engineering and providing encouragement and support to the creative thinking process. We envision a system where users are able to refine their questions, explore topics in greater depth, and iteratively develop original ideas with the assistance of AI. Future versions of the system may expand this database to include a broader range of open-access sources, such as scholarly paper websites and digital libraries. Furthermore, case studies may be conducted among small groups of high school or university students. Participants will have the opportunity to interact with the system for their brainstorming stages of a creative project. Qualitative data can then be gathered from their responses of the system's performance and the level of creativity and originality compared to other students using generative AI and/or completing the same creativity tasks on their own.

Acknowledgements

The work is supported by National Science Foundation (#2142783). We thank all reviewers for the valuable feedback.

References

Anderson, B. R., Jash Hemant Shah, & Kreminski, M. 2024. Homogenization Effects of Large Language Models on Human Creative Ideation. *Creativity and Cognition*, 24. https://doi.org/10.1145/3635636.3656204

Boussioux, L., Lane, J. N., Zhang, M., Jacimovic, V., & Lakhani, K. R. 2024. The Crowdless Future? Generative AI and Creative Problem-Solving. *Organization Science*, 35(5). https://doi.org/10.1287/orsc.2023.18430

Cropley, D. H. 2025. "The Cat Sat on the …?" Why Generative AI Has Limited Creativity. *The Journal of Creative Behavior*, 59(4). https://doi.org/10.1002/jocb.70077

Fatemi, G., & Saito, E. 2019. Unintentional plagiarism and academic integrity: The challenges and needs of postgraduate international students in Australia. *Journal of Further and Higher Education*, 44(10), 1–15. https://doi.org/10.1080/0309877x.2019.1683521

Habib, S., Vogel, T., Thorne, E., & Xiao, A. 2023. How Does Generative Artificial Intelligence Impact Student Creativity? *Journal of Creativity*, 34(1). https://doi.org/10.1016/j.yjoc.2023.100072

Holmström, J., & Carroll, N. 2024. How organizations can innovate with generative AI. *Business Horizons*. https://doi.org/10.1016/j.bushor.2024.02.010

Panda, S., & Kaur, D. N. 2024. Exploring the Role of Generative AI in Academia: Opportunities and Challenges. Ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4915738

Rafner, J., Beaty, R. E., Kaufman, J. C., Lubart, T., & Sherson, J. 2023. Creativity in the age of generative AI. *Nature Human Behaviour*, 7(11), 1836–1838. https://doi.org/10.1038/s41562-023-01751-1

Stevenson, C., Smal, I., Baas, M., Grasman, R., & van der Maas, H. 2022. Putting GPT-3's Creativity to the (Alternative Uses) Test. arXiv:2206.08932.

Tan, C., Chuah, C, Lee, S, Tan, C. 2021.Being Creative Makes You Happier: The Positive Effect of Creativity on Subjective Well-Being. *International Journal of Environmental Research and Public Health*. 18(14) 7244. https://doi.org/10.3390/ijerph18147244.

Yu, L., Alégroth, E., Chatzipetrou, P., & Gorschek, T. 2025. Evaluating the quality of GenAI applications in software engineering: a multi-case study. *Empirical Software Engineering*, 31(2). https://doi.org/10.1007/s10664-025-10759-2.

Zhou, E., & Lee, D. 2024. Generative artificial intelligence, human creativity, and art. *PNAS Nexus*, 3(3). https://doi.org/10.1093/pnasnexus/pgae052.